

Thermal Aware Scheduling on FPGAs

Yingyi Luo, Xiaoyang Wang,
Gokhan Memik, Seda Ogrenci-Memik,
Kazutomo Yoshii and Pete Beckman

Sep 13, 2017

Northwestern

Agenda


- Introduction
- Motivation
- Model
- Experiment
- Conclusion

Agenda

- Introduction
- Motivation
- Model
- Experiment
- Conclusion

Ideal Case vs Actual Case

Ideally




Actually





Our Work on FPGAs


Uneven distribution



Even distribution



Cool
them all!



Zzz...

Agenda


- Introduction
- **Motivation**
- Model
- Experiment
- Conclusion

Chameleon Infrastructure

- 2U servers
- Nallatech 385A with Arria 10 chip
- Intel FPGA examples (written in OpenCL)
- Located vertically on the same rack
- Positioning numbers start at the bottom



U47_48



U45_46




U43_44




U41_42

Thermal Variation in HPC

Peak temperature variation across FPGA boards




A Quick Example




A

57.81




B

80.82




C

59.84




D

57.14




C

61.87




D

63.23



B

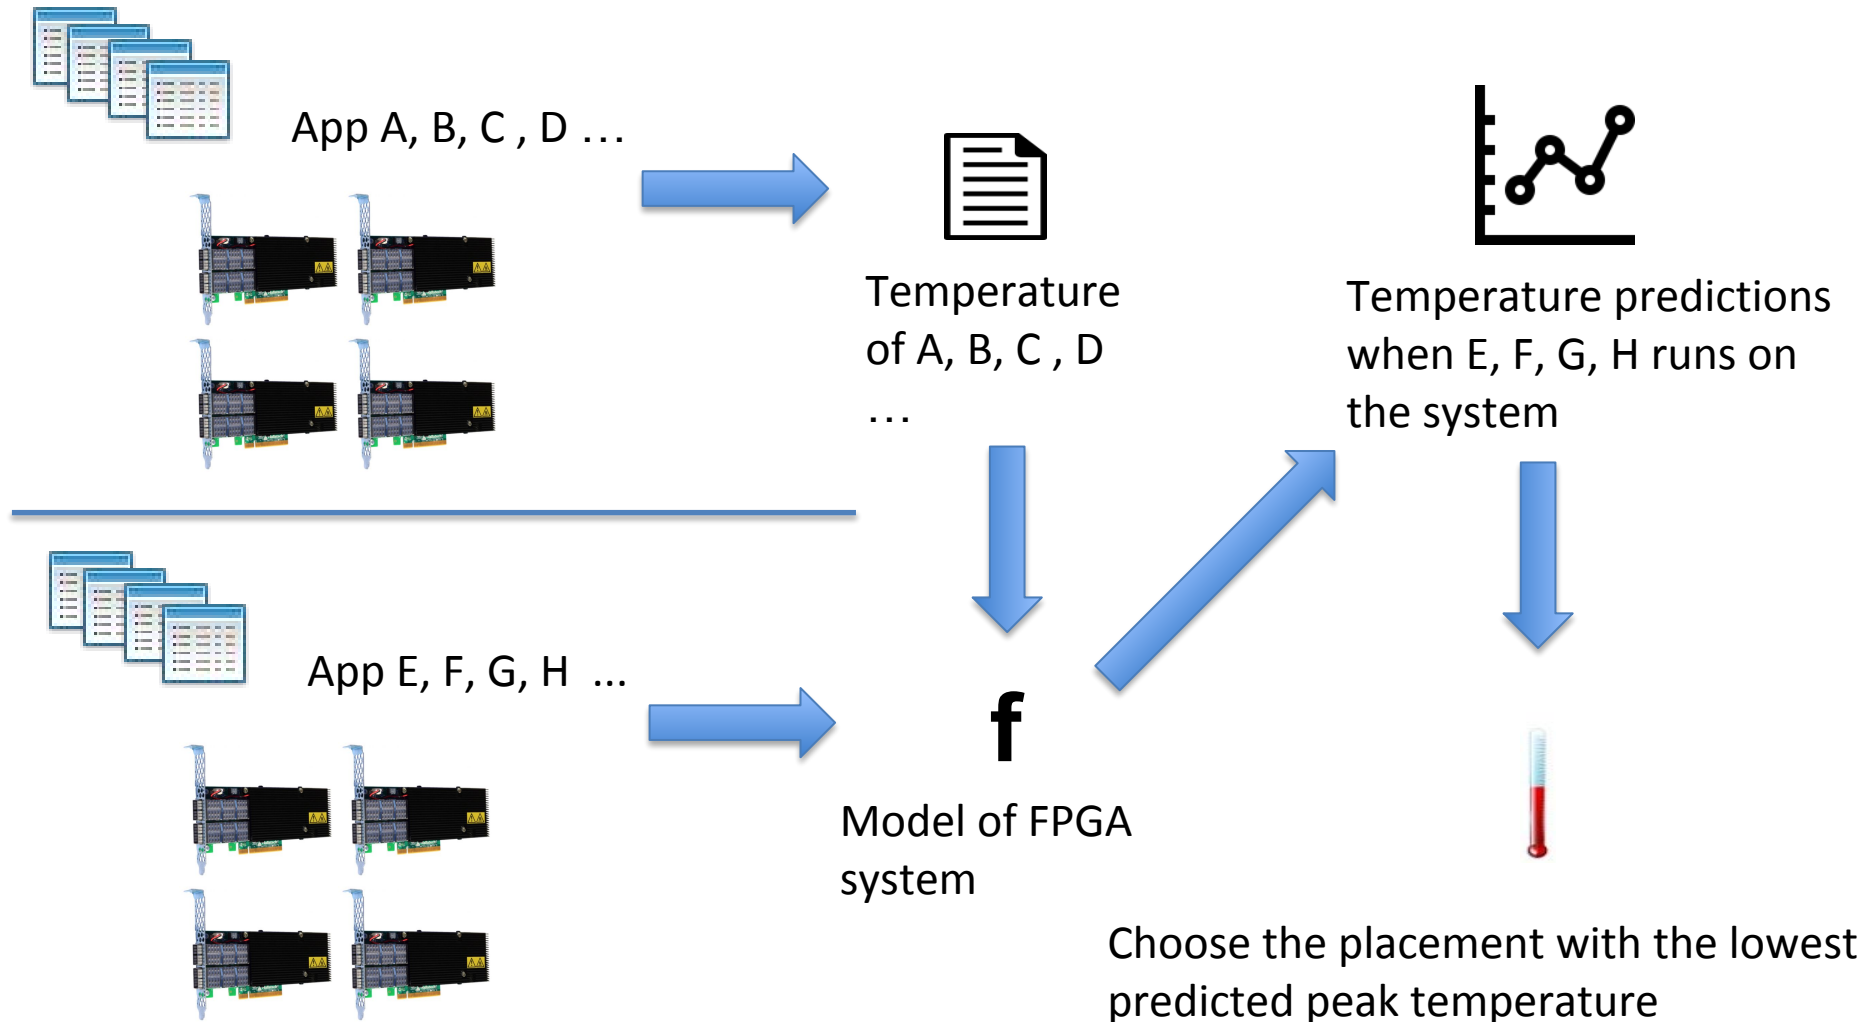
69.32



A


58.49

Bird's Eye View




Leakage Power and Temperature

Peak temperature and peak power relation



Consistent Performance

Normalized FPGA performance across machines



Agenda

- Introduction
- Motivation
- **Model**
- Experiment
- Conclusion

Machine Learning Model

Learner

Linear Regression

Neural Network

Nearest Neighbor

Random Forest

Label

Peak Temperature

Feature

Logic Utilization

RAM Blocks

Frequency

DSP Blocks

Memory Bits

I/O Pins

The Training Sample

- 210 combinations, select 4 tasks from 10 tasks
- 24 placements, place 4 tasks on 4 machines. One placement is one sample

5040 samples, in total

- 12 features, each task has 3 features
- 1 label, each placements has 1 peak temperature

12 features and 1 label, each sample

Prediction Model

Task Dependent Model

Split the samples into 2 sets, one for training(80%), the other one for testing(20%)

Task Independent Model

Build the model with 6 tasks and use the remaining 4 to test it

The result will come in Experiment section


Scheduling Methodology

- Take the task dependent model as example
- Build a prediction model with 4056 training samples (80%)
- When a task combination comes (from the remaining samples), enumerate 24 placements
- Make prediction for each placements
- Choose the lowest predicted peak temperature


Agenda

- Introduction
- Motivation
- Model
- **Experiment**
- Conclusion

Prediction Statistics




Scheduler Performance



The average peak temperature of the system when we use the schedulers

The advantage we expect to get from our schedulers




Power Reduction

Peak Temperature Reduction: **4.60 °C**


Peak Power Reduction: **1.77 W**

Power Sum Reduction: **1.97 W**




Scheduler Performance

Task Independent Model




A Closer Look, Linear Regression

Task Independent Model



Scheduler Performance

Task Independent Model



Agenda

- Introduction
- Motivation
- Model
- Experiment
- Conclusion

Conclusion

- Up to 11.50°C peak temperature variation across machines for the same benchmark
- First study in this area to our best knowledge
- 4.36°C peak temperature reduction on average
- Build machine learning models and develop schedule methodology to explore potential improvement